

GloTok: Global Perspective Tokenizer for Image Reconstruction and Generation

Xuan Zhao^{1*†}, Zhongyu Zhang^{2*}, Yuge Huang², Yuxi Mi¹, Guodong Mu², Shouhong Ding^{2‡}, Jun Wang³, Rizen Guo³, Shuigeng Zhou^{1‡}

¹College of Computer Science and Artificial Intelligence, Fudan University, Shanghai, China

²Tencent Youtu Lab, Shanghai, China

³Tencent WeChat Pay Lab33, Shanghai, China

xzhao23@m.fudan.edu.com, {yxmi20, sgzhou}@fudan.edu.com

{wilxyzhang, yugehuang, gordonmu, ericshding}@tencent.com

{earljwang, rizenguo}@tencent.com

Abstract

Existing state-of-the-art image tokenization methods leverage diverse semantic features from pre-trained vision models for additional supervision, to expand the distribution of latent representations and thereby improve the quality of image reconstruction and generation. These methods employ a locally supervised approach for semantic supervision, which limits the uniformity of semantic distribution. However, VAE proves that a more uniform feature distribution yields better generation performance. In this work, we introduce a Global Perspective Tokenizer (GloTok), which utilizes global relational information to model a more uniform semantic distribution of tokenized features. Specifically, a codebook-wise histogram relation learning method is proposed to transfer the semantics, which are modeled by pre-trained models on the entire dataset, to the semantic codebook. Then, we design a residual learning module that recovers the fine-grained details to minimize the reconstruction error caused by quantization. Through the above design, GloTok delivers more uniformly distributed semantic latent representations, which facilitates the training of autoregressive (AR) models for generating high-quality images without requiring direct access to pre-trained models during the training process. Experiments on the standard ImageNet-1k benchmark clearly show that our proposed method achieves state-of-the-art reconstruction performance and generation quality.

1 Introduction

Recent years have witnessed remarkable advancements in image generation, driven by autoregressive (AR) models (Vaswani et al. 2017; Radford et al. 2018; Touvron et al. 2023) and diffusion models (Dhariwal and Nichol 2021; Ho and Salimans 2022; Rombach et al. 2022). To promote high-quality generation, many contemporary methods adopt a two-stage paradigm: (1) a tokenizer compresses input images into tokens, followed by (2) a generator modeling the

*These authors contributed equally.

†This work was done by Xuan Zhao during an internship at Tencent Youtu Lab.

‡Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

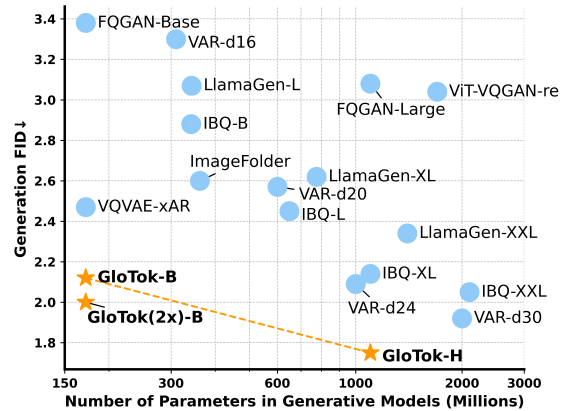


Figure 1: Comparison of generative performance among different methods and GloTok, where lower values on the Y-axis correspond to better performance.

distribution of tokens. The compressed tokens determine both structural composition and detailed characteristics of generated images. Therefore, image tokenizers play a crucial role in improving the performance of image generation.

Recent SOTA tokenizers (Zhu et al. 2024c; Qu et al. 2024; Li et al. 2024b; Bai et al. 2024; Ma et al. 2025a) for image generation with AR models leverage pre-trained models such as DINO (Caron et al. 2021; Oquab et al. 2023; Darcet et al. 2023) and CLIP (Radford et al. 2021) as additional supervision. Concretely, the pre-trained model converts the input images into semantic features, and these features are then served to guide the tokenizer to generate more effective latent representations for image generation.

While these approaches have achieved remarkable performance in enhancing the information capacity of latent representations, their dependence on single-image semantic supervision poses a significant challenge to semantic codebook fitting. Specifically, this local approach is confined to performing contrastive learning between latent-space features and pre-trained features within individual images, thereby

constraining its capacity to model feature discrepancies across the entire dataset. Consequently, the learned latent-space distribution exhibits suboptimal breadth and uniformity. VA-VAE (Yao, Yang, and Wang 2025) demonstrates that more uniform latent representations can significantly improve generative performance. Based on the conclusions of this prior work, suboptimal uniform feature distributions can tend to constrain generative performance.

In this paper, we introduce GloTok, a novel tokenizer that outperforms previous work in Fig. 1 and is designed to achieve a more uniform latent distribution, leveraging global relations across the entire dataset. In practice, GloTok adopts a dual codebook architecture to learn the semantics and details of the input images. A codebook-wise histogram relation learning method is proposed to transfer the semantics to the semantic codebook by feature relationships modeled by pre-trained models across the entire dataset (called global relations). Taking ImageNet as an example, existing methods require batch-based semantic learning for 1.28 million images, resulting in unstable optimization and makes it difficult to achieve global optimality. Differently, our approach consistently leverages global relations as the semantic guidance, such as a multi-bin histogram to summarize the global relations. This design not only facilitates improved uniformity of semantic latent distribution, but also eliminates the need to directly access the pre-trained model during training, thereby reducing both temporal and spatial overhead.

In addition, to mitigate accuracy degradation caused by discretization, we employ a residual learning module in GloTok. This module takes quantized features as input, predicts the residual between the original continuous features and their quantized discrete counterparts, and fuses this residual with the quantized features to produce the final quantized feature representations. Residual learning enables the tokenizer to partially restore fine-grained details, thereby improving the quality of image reconstruction.

Through a novel histogram relation learning method and residual learning modules, the performance of the AR model trained with our proposed GloTok achieves significant improvements in generative tasks. Our experiments demonstrate that the proposed model achieves a state-of-the-art (SOTA) reconstruction FID of 0.83 at 256×256 resolution on the standard ImageNet benchmark. And we get a competitive generation FID=1.75 on the ImageNet 256×256 image generation using an autoregressive generator.

Our contribution can be summarized as follows.

- We propose GloTok, a novel tokenizer that constrains the global relational distribution within the latent space to transfer the semantics from pre-trained models, achieving a significant improvement in the performance of image reconstruction and generation.
- A codebook-wise histogram relation learning method is introduced to transfer the semantics to the semantic codebook, which models a more uniform latent distribution and eliminates the need to directly access the pre-trained models during training.
- To mitigate accuracy degradation induced by discretization, a residual learning module is designed to predict and

fuse token residuals into the final representations, preserving fine-grained details for superior image quality.

2 Related Work

2.1 Image Tokenization

Image tokenization is the fundamental process in image reconstruction and generation. An image tokenizer converts the input image into continuous or discrete tokens.

As a pioneering work on transformer-based image generation, VQ-GAN (Esser, Rombach, and Ommer 2021) proposes the encoder-quantizer-decoder architecture, which employs a vector quantization (VQ) technique to map the continuous latent space features of the input images into discrete tokens. However, the performance of VQGAN is constrained by both a small codebook size and a low token utilization rate. ViT-VQGAN (Yu et al. 2021) and Efficient-VQGAN (Cao et al. 2023) leverage the Vision Transformer (ViT) architecture (Dosovitskiy et al. 2020), which mitigates reliance on convolutional inductive biases due to its global attention mechanisms. RQ-VAE (Lee et al. 2022) adopts a residual quantization (RQ) to precisely approximate the feature map. VQGAN-LC (Zhu et al. 2024a) and SimVQ (Zhu et al. 2024b) employ linear layers to expand the codebook capacity while preserving high codebook utilization efficiency. Furthermore, MAGViT-v2 (Yu et al. 2023) proposes a lookup-free quantization method leading to a large vocabulary. Unlike typical 2D tokenizers, TiTok (Yu et al. 2024b) adopts a 1D tokenization framework that enables a more flexible token count design while also capturing semantically rich image information.

Recent advancements have focused on injecting semantic features from pre-trained vision models such as DINO (Caron et al. 2021; Oquab et al. 2023; Darcet et al. 2023) or CLIP (Radford et al. 2021) into the codebook to enhance the generation performance of VQGAN. DiGiT (Zhu et al. 2024c) integrates the discrete tokens derived from clustered DINO (Oquab et al. 2023) features into the latent space of VQGAN to enrich the semantic information of the latent space of the generative model. UniTok (Ma et al. 2025a) adopts a pre-trained CLIP (Radford et al. 2021) model to align the latent space of input images with the semantic latent space derived from the corresponding text. Recent work, such as ImageFolder (Li et al. 2024b) and FQGAN (Bai et al. 2024), utilizes a multi-codebook architecture in which certain codebooks are designed to capture the feature distributions extracted from pre-trained vision models using a feature-wise contrastive loss. These methods adopt a local perspective, specifically by capturing semantics through intra-image feature comparison with pre-trained models. In contrast, our tokenizer GloTok employs a method that globally constrains the distribution of semantic features with pre-trained relational distributions to transfer the semantics.

2.2 AR Image Generation

Recent years have witnessed remarkable advancements in auto-regressive (AR) image generation, where models synthesize images by capturing the distribution of discrete tokens. Early foundational works, such as VQGAN (Esser,

Rombach, and Ommer 2021), VIT-VQGAN (Yu et al. 2021), and RQ-Transformer (Lee et al. 2022), employ transformer architecture to model discrete token sequences for image generation. Building on this, LlamaGen (Sun et al. 2024) introduces the Llama (Touvron et al. 2023) model to enhance image generation capabilities. Beyond architectural innovations, optimizing token training strategies has emerged as a key research direction. For instance, VAR (Tian et al. 2024) adopts a multi-resolution token prediction framework to improve the image fidelity and detail preservation. MAR (Li et al. 2024a) models the per-token probability distribution using a diffusion procedure, which achieves strong generation results. FlowAR (Ren et al. 2024) introduces a general next-scale prediction method with a streamlined scale design. RAR (Yu et al. 2024a) randomly permutes the input token sequence into different factorization orders for modeling bidirectional contexts. xAR (Ren et al. 2025) enables flexible prediction units and effectively alleviates exposure bias, achieving significant generative performance.

3 Method

3.1 Preliminary

AR image generation models are trained from discrete tokens generated from an image tokenizer. The tokenizer typically follows an encoder-quantizer-decoder architecture, where the encoder extracts features from the image, the quantizer converts these features into discrete tokens, and the decoder reconstructs the image from these tokens. This process allows the model to effectively learn and generate images based on tokenized representations.

Given an image $x \in \mathcal{R}^{H \times W \times 3}$, the encoder E first extracts continuous features $Z^{h \times w \times c}$ from the input image x , where $h \times w$ represents the sequence length of the latent tokens and c represents the channel dimension. For image quantization, the quantizer selects the closest token t for each feature z of Z with a learnable codebook $C^{n \times d}$ by calculating the Euclidean distance (Esser, Rombach, and Ommer 2021) or the cosine similarity (Shi et al. 2024; Ma et al. 2025b; Bai et al. 2024). The quantization loss L_Q is calculated based on the quantized features $\hat{Z}^{h \times w \times c}$ and the original features Z to optimize the latent space and the codebook:

$$L_Q = \left\| \text{sg}[Z] - \hat{Z} \right\|_2^2 + \beta \left\| \text{sg}[\hat{Z}] - Z \right\|_2^2 \quad (1)$$

Here, $\text{sg}[\cdot]$ denotes the stop-gradient operation and $\|\cdot\|_2^2$ denotes the Euclidean distance (or called L2 distance). The quantization process optimizes latent representations within the codebook while reducing the representation gap between the continuous latent space and the discrete latent space.

After quantization, the decoder D decodes \hat{Z} into the reconstruction image \hat{x} . For the representative VQGAN (Esser, Rombach, and Ommer 2021) model, the optimization of the entire network is represented by the following loss function L_{VQGAN} :

$$L_R = \|x - \hat{x}\|^2 \quad (2)$$

$$L_{VQGAN} = \lambda_R L_R + \lambda_Q L_Q + \lambda_P L_P + \lambda_G L_G \quad (3)$$

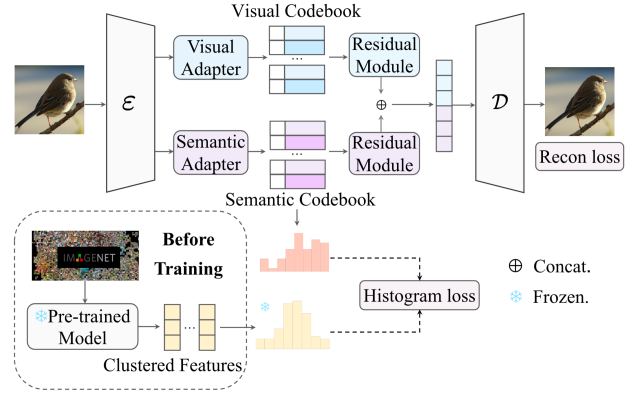


Figure 2: Illustration of our method. Top: GloTok encoder-quantizer-decoder architecture with dual codebooks and residual modules. Bottom: overview of the Histogram Relation Learning method. An image is quantized into two sets of tokens by a visual codebook and a semantic codebook. The semantic codebook learns the token relationship from features clustered from a pre-trained model with a histogram loss. GloTok adopts two residual modules to learn the residuals between continuous features and discrete features.

where λ_i is the weight factor for each loss term, L_R is the reconstruction loss, L_P is the perceptual loss between x and \hat{x} , and L_G is adversarial loss encouraging to learn more realistic images x with a patch-discriminator (Isola et al. 2017).

3.2 Architecture

Similarly to previous methods (Bai et al. 2024; Li et al. 2024b), as shown in Fig. 2, we leverage a dual-codebook architecture with a convolutional encoder and decoder. we encode the input image into two sets of tokens: one consisting of the semantic tokens and the other of the visual tokens. We use the semantic tokens to model the distribution of semantic features, while visual tokens are employed to model the image’s visual features. Specifically, the encoder E first encodes the input image x into the latent features $Z^{h \times w \times c}$. We then adopt two adapters to decompose Z into semantic features Z_{sem} and visual features Z_{vis} . For quantization, Z_{sem} and Z_{vis} are quantized into \hat{Z}_{sem} and \hat{Z}_{vis} respectively, with learnable codebooks C_{sem} and C_{vis} .

To control the global feature space distribution of semantic features, we introduce a novel codebook-wise histogram relation learning method that constrains the distribution of relations between tokens in the semantic codebook, detailed in Section 3.3. The visual codebook is specifically served to capture visual feature representations within the latent space, enabling image reconstruction by leveraging structural patterns learned from the semantic codebook.

Quantization of the continuous latent Z into discrete representations inherently induces information degradation. To mitigate the limitation, we designed a residual learning module to recover lost information, thus improving the quality of image reconstruction. The learned residual features are

aggregated with quantized tokens through summation operations to generate information-complete token representations. Subsequently, the decoder decodes the composite latent representation \hat{Z} formed by channel-wise concatenation of \hat{Z}_{sem} and \hat{Z}_{vis} to synthesize the output image.

3.3 GloTok

Histogram relation learning In order to visualize the distribution of the quantitative data, the histogram is a commonly used method that reflects the probability density of the underlying data distribution. Inspired by this concept, we propose a codebook-wise histogram relation learning method that transfers the prior global relational distribution knowledge to the semantic codebook.

To achieve this objective, it is necessary to select an appropriate source for the prior distribution. We adopt two approaches for its construction: (1) directly utilizing the semantic codebook of the tokenizer from a pre-trained model; (2) performing k-means clustering on the intermediate layer features of CLIP (Radford et al. 2021) or DINOv2 (Oquab et al. 2023) adopted in VQGAN-LC (Zhu et al. 2024a).

In particular, our methodology differs from VQGAN-LC in that it employs an MLP layer to establish the mapping from clustered features to the codebook. Although the feature sources in VQGAN-LC contain rich semantic information, the semantic constraints imposed on the mapped codebook are insufficient, and the use of an MLP layer for training a large codebook, which is up to 100,000, results in substantial computational cost. Instead, we focus on the patterns of relational similarity between tokens.

To construct the global relational distribution, we first compute the pairwise cosine similarities between tokens within each token set $C^{n \times d}$:

$$D = \{f(t_i, t_j) | t_i, t_j \in C, i \leq j\} \quad (4)$$

where $f[\cdot]$ is a *cosine similarity* function. Subsequently, in a histogram-like manner, we initialize N bins slicing from -1 to 1, where each bin corresponds to a specific sub-range of quantized token distances:

$$B = \left\{ -1 + \frac{2k}{N-1} \mid k = 0, 1, 2, \dots, N-1 \right\} \quad (5)$$

To quantify the distribution of D across different intervals B , we compute a Gaussian-smoothed histogram of pairwise cosine similarities:

$$p_n = \sum_{i=1}^{\|D\|} \exp(-\alpha(D_i - B_n)^2) \quad (6)$$

where $\|D\|$ is the dim of D and α is a weight factor.

We define the relational distribution of the trainable codebook as the student distribution P_s and define the relational distribution of pre-trained tokens as the teacher distribution P_t . Then, we employ a Kullback-Leibler divergence to constrain the student distribution with the teacher distribution:

$$L_{hist} = KL(\log P_s \| sg[P_t]) \quad (7)$$

where $sg[\cdot]$ denotes the stop-gradient operation, P_s denotes the normalized student distribution, and P_t denotes the normalized teacher distribution.

Residual learning The quantization process, which transforms continuous features into discrete representations, inherently induces precision degradation. While discrete representations simplify the latent space expression, they inevitably compromise the model’s representational capacity.

We propose a residual learning method in which residual modules R_{sem} and R_{vis} are implemented as transformer blocks process quantized latent features \hat{Z}_{sem} and \hat{Z}_{vis} , respectively, to predict lost latent details during discretization:

$$Z_{res,sem} = R_{sem}(\hat{Z}_{sem}) \quad (8)$$

$$Z_{res,vis} = R_{vis}(\hat{Z}_{vis}) \quad (9)$$

The optimization is derived from the differential mapping between continuous features Z and the quantized features \hat{Z} through residual error minimization:

$$L_{res} = \left\| sg[Z_{sem}] - \hat{Z}_{sem} - Z_{res,sem} \right\| + \left\| sg[Z_{vis}] - \hat{Z}_{vis} - Z_{res,vis} \right\| \quad (10)$$

Discrete features \hat{Z}_{sem} and \hat{Z}_{vis} are combined with their respective learned residuals through an element-wise addition operation to form the final discrete representation:

$$\hat{Z}_{vis} = \hat{Z}_{vis} + Z_{res,vis} \quad (11)$$

$$\hat{Z}_{sem} = \hat{Z}_{sem} + Z_{res,sem} \quad (12)$$

Training loss Our model is trained with a combined loss:

$$L = L_{VQGAN} + \lambda_{res}L_{res} + \lambda_{hist}L_{hist} \quad (13)$$

where λ_{res} and λ_{hist} are weight factor that respectively control L_{res} and L_{hist} , thus balancing various loss terms.

3.4 Training with Auto-Regressive Model

Through quantization, the input image is discretized into two complementary token sets in our tokenizer. Previous auto-regressive models (Esser, Rombach, and Ommer 2021; Touvron et al. 2023; Tian et al. 2024) are trained with the indices of the tokens quantized by codebook. These auto-regressive models are designed to model inter-token dependencies via sequential pattern learning.

However, although previous multi-branch models (Bai et al. 2024; Li et al. 2024b) exhibited superior reconstruction performance compared to other single-branch methods (Sun et al. 2024; Tian et al. 2024), their generation performance remained relatively suboptimal. When the tokenizer employs a multi-branch architecture, it must generate multiple indices to represent images. Due to the exposure bias (Arora et al. 2023) inherent in AR training, producing multiplicative indices exacerbates index prediction errors during inference, thus degrading generation performance. In our work, we aim to circumvent this limitation.

We concatenate the semantic features and visual features obtained after vector quantization along the channel dimension, and employ an xAR (Ren et al. 2025) model to directly learn this combined feature space. This approach enhances the generation performance of the multi-branch tokenizer.

Codebook Type	Method	Codebook Size	rFID ↓	PSNR ↑
Single	VQGAN (Esser, Rombach, and Ommer 2021)	16,384	5.96	23.3
Single	RQ-VAE (Lee et al. 2022)	16,384	3.20	-
Single	VQGAN-LC (Zhu et al. 2024a)	16,384	3.01	23.2
Single	VQGAN-LC (Zhu et al. 2024a)	100,000	2.62	23.8
Single	LlamaGen (Sun et al. 2024)	16,384	2.19	20.7
Single	MaskBit (Weber et al. 2024)	16,384	1.37	21.5
Single	TiTok-S-128 (Yu et al. 2024b)	4,096	1.71	17.7
Single	Open-MAGVIT2 (Luo et al. 2024)	16,384	1.58	-
Single	Open-MAGVIT2 (Luo et al. 2024)	262,144	1.17	22.6
Single	IBQ (Shi et al. 2024)	16,384	1.37	21.0
Single	IBQ (Shi et al. 2024)	262,144	1.00	-
Muti	TokenFlow (Qu et al. 2024)	32,768 × 2	1.37	21.4
Muti	FQGAN-Dual (Bai et al. 2024)	16,384 × 2	0.94	22.0
Muti	GloTok	16,384	0.92	22.4
Muti	GloTok(2x)	16,384 × 2	0.83	22.7

Table 1: Reconstruction comparison with other tokenizers evaluated on ImageNet 256×256 50k validation dataset.

4 Experiments

4.1 Experiments Setting

Implementation details We trained and evaluated GloTok on the 256×256 ImageNet (Deng et al. 2009) benchmark. The encoder and decoder settings are consistent with VQGAN (Esser, Rombach, and Ommer 2021). To facilitate alignment in methods employing single-branch codebook and multi-branch codebook settings with the same codebook size, respectively, we trained two tokenizers with distinct codebook size configurations: (1) the semantic codebook is configured with a size of 4,096, and the visual codebook with a size of 12,288, resulting in a total size of 16,384; (2) both the semantic and visual codebooks have a size of 16,384, which we refer as GloTok(2x). The dimension of the codebooks is 8. To transfer inter-token relational knowledge from pre-trained models, we train the tokenizer with diverse teacher distributions, such as K-means-clustered DINOv2 tokens and a pre-trained semantic codebook. We configure the number of bins (B) of the histogram learning method to 40. We set $\lambda_{hist} = 0.01$, $\lambda_{res} = 0.5$. Within the residual learning module, we incorporate a single-layer transformer block to predict residual components. The tokenizer is trained with the following settings: a learning rate of $2e-4$ without any decay mechanism, an Adam Optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.95$, a global batch size of 240 with 150 epochs. For the generator, we train FAR-B, FAR-L with 200 epochs with our tokenizer, and we train xAR-B, xAR-H (Ren et al. 2025) with 800 epochs following xAR settings.

Metrics We employ the Frchet Inception Distance (FID) (Heusel et al. 2017) and the Peak Signal-to-Noise Ratio (PSNR) to evaluate reconstruction quality. We use the ImageNet validation set, consisting of 50k samples, to compute the reconstruction FID (rFID) and PSNR. To evaluate the quality of generated samples, we employ two widely recognized metrics: generation FID (gFID) and Inception Score (IS) (Salimans et al. 2016). The gFID measures the similarity between the distribution of generated images and that of real-world data, providing a quantitative assessment of the

realism and diversity of the outputs. On the other hand, the IS evaluates both the quality and diversity of the generated samples by leveraging an Inception model pretrained on the ImageNet dataset.

4.2 Main Results

Image reconstruction As shown in Tab. 1, our tokenizer was compared with other state-of-the-art models in terms of reconstruction performance. Reconstruction performance of our tokenizer was assessed at 256×256 resolution, employing 16x downsampling, on the ImageNet-1K validation set.

Our tokenizer achieves state-of-the-art (SOTA) reconstruction performance on latent resolution 16×16 , attaining a reconstruction FID of 0.83 at 16384×2 codebook size. Notably, the multi-codebook architecture demonstrates substantially superior reconstruction FID compared to single-codebook architectures, validating the effectiveness of learning hierarchical latent space representations across different aspects. Furthermore, our model at a total 16384 codebook size, attaining an rFID=0.92, not only outperforms tokenizers with equivalent codebook size, but also achieves superior performance relative to those with larger codebooks.

Image generation We present a comprehensive summary of the experimental results on the ImageNet-1K image generation benchmark at a resolution of 256×256 pixels in Tab. 2. Given that the xAR method directly learns the quantized latent representations rather than relying on indices, we conducted an evaluation to assess the effectiveness of the quantized features of GloTok. We trained xAR using GloTok configured with a total codebook size of 16384, achieving a competitive generative performance and attaining a gFID score of 1.75 and an IS of 327.50. It should be noted that xAR gives the generation results in the VAE (Kingma and Welling 2022) feature space without quantization. In order to make a fair comparison with the quantization feature space, we trained xAR with a VQ-VAE-based tokenizer and achieved a gFID score of 2.47. In comparison, when we trained xAR using GloTok with a total codebook size

Type	Method	Params	Tokens	gFID↓	IS↑
AR	VQGAN (Esser, Rombach, and Ommer 2021)	227M	256	18.65	80.4
AR	VQGAN-LC (Zhu et al. 2024a)	404M	256	15.4	-
AR	LlamaGen-B (Sun et al. 2024)	111M	256	5.46	193.61
AR	LlamaGen-XXL (Sun et al. 2024)	1.4B	256	3.09	253.60
AR	LlamaGen-3B (Sun et al. 2024)	3.1B	256	3.06	279.71
AR	IBQ-B (Shi et al. 2024)	342M	256	2.88	254.73
AR	IBQ-XXL (Shi et al. 2024)	2.1B	256	2.05	286.73
VAR	ImageFolder (Li et al. 2024b)	362M	256	2.60	295.0
VAR	VAR-d16 (Tian et al. 2024)	310M	256	3.30	274.4
VAR	VAR-d24 (Tian et al. 2024)	1.0B	256	2.09	312.9
VAR	VAR-d30 (Tian et al. 2024)	2.0B	256	1.92	323.1
FAR	FQGAN (Bai et al. 2024)	415M	512	3.38	248.26
FAR	FQGAN (Bai et al. 2024)	898M	512	3.08	272.52
FAR	GloTok	415M	512	3.13	280.66
FAR	GloTok(2x)	415M	512	2.95	261.38
FAR	GloTok	898M	512	2.98	285.34
xAR	VQ-VAE†	172M	256	2.47	277.63
xAR	GloTok	172M	256	2.12	272.16
xAR	GloTok(2x)	172M	256	2.00	269.25
xAR	GloTok	1.1B	256	1.75	327.50

Table 2: Class-conditional generation comparison on ImageNet-1k 256×256 resolution. The evaluation protocol and implementation are the same as ADM. †denotes that the xAR is trained with a VQ-VAE tokenizer from LlamaGen (Touvron et al. 2023).



Figure 3: ImageNet-1k 256×256 generated samples of GloTok trained with xAR.

of 16384, the resulting gFID score dropped to 2.12, a 0.35 reduction relative to the VQ-VAE-based baseline. Fig. 3 shows the visualization results of our 16384 codebook size tokenizer’s 256×256 generated images training with xAR, showcasing its high fidelity. In order to validate the impact of histogram relation learning on the generative performance, we trained the FAR-B model using a tokenizer with the same codebook size (16384×2) as that of the FQGAN-dual model, which employs a local contrastive learning approach.

Compared to FQGAN, our method achieved a gFID reduction of 0.43 and an IS improvement of 13, demonstrating that our approach yields superior generative performance relative to local contrastive methods. These findings further validate the effectiveness of the histogram relation learning method in enhancing generative performance.

4.3 Histogram Relation Learning Makes More Uniform

We select quantized semantic latent representations for evaluation, like VA-VAE (Yao, Yang, and Wang 2025). We randomly select images from the Imagenet-1k val dataset. We then get the semantic latent representations \hat{Z}_{sem} from FQGAN and GloTok with the same 16384×2 codebook size. For VQ-VAE, we select a tokenizer from LlamaGen (Touvron et al. 2023) with a 16384 codebook size as a baseline. The dimension of tokens from the above codebooks is all 8. We computed the standard deviation and Gini coefficients to evaluate the uniformity. Tab. 3 shows that the semantic latent distribution of GloTok is more uniform than FQGAN using a pre-trained model for local contrastive learning.

As shown in Tab. 2, the FAR trained with GloTok utilizing a codebook size of 16384×2 attained a gFID score of 2.95, which surpasses the gFID score of 3.38 achieved by the FAR-B trained with FQGAN. This finding demonstrates that a more uniform semantic space distribution is more beneficial for improving generative performance.

4.4 Ablation Studies

To validate the effectiveness of our proposed framework, we conducted ablation studies on GloTok with a 4096 semantic codebook size and a 12288 visual codebook size.

Tokenizer	density cv↓	normalized entropy↑	gini coefficient↓
VQ-VAE†	0.221	0.9971	0.1191
FQGAN	0.225	0.9970	0.1202
GloTok	0.217	0.9972	0.1168

Table 3: Comparison of uniformity metrics of VQ-VAE, FQGAN, GloTok. †denotes that the VQ-VAE is obtained from LlamaGen (Touvron et al. 2023).

The impact of different feature levels As illustrated in Fig. 4, to evaluate the impact of different feature levels on overall image reconstruction, we visualize how varying retention ratios of visual features (ranging from partial to full preservation) influence image quality both with and without semantic feature guidance. Specifically, we set the discarded features to 0. We retain 100% of the semantic features while varying the retention ratio of the visual features from 25% to 100% to observe their interaction.

The results show that visual features alone preserve substantial color information in reconstructed images, while structural and textural details are diminished. However, when semantic information is incorporated, the structural and textural components are effectively restored. The difference indicates that the semantic codebook integrates high-frequency components (e.g., structural and textural elements), while the visual codebook combines low-frequency components (e.g., color distributions). Our experiments demonstrate the effectiveness of the global relational learning method in transferring semantics into the codebook.

The effect of the histogram relation learning Our tokenizer is trained for 40 epochs under varying weight configurations of histogram relation learning to investigate its impact on the reconstruction performance of our tokenizer. As shown in Tab. 4, our experimental analysis reveals that incorporating the method facilitates a reduction in rFID from 1.43 to 1.21, accompanied by an improvement in PSNR from 22.06 to 22.23. These findings demonstrate that histogram relation learning plays a pivotal role in enhancing tokenizer reconstruction performance. Empirically, setting the hyperparameter weight to 0.01 achieved superior optimization outcomes in our experiments.

Histogram Relation	Weight	rFID ↓	PSNR ↑
×	-	1.43	22.06
✓	0.1	1.27	21.97
✓	0.01	1.21	22.23

Table 4: Effect of the histogram relation learning. Tokenizers are only trained by the histogram relation learning method without the residual module.

Different components We evaluated the contributions of individual components to the improvement of GloTok reconstruction. Results are reported in Tab. 5. We start with a vanilla dual-codebook design. Through the adoption of

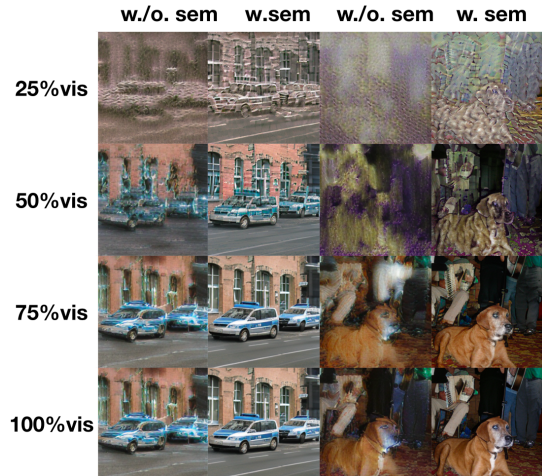


Figure 4: Visualization of the reconstruction on different retention rates of visual features. Each row presents visualization results for visual features with different retention ratios, both with and without semantic features.

histogram relation learning, our experiments demonstrate a significant improvement in reconstruction quality, with the Fréchet Inception Distance (FID) decreasing from 1.43 to 0.97 and the Peak Signal-to-Noise Ratio (PSNR) increasing from 22.06 to 22.30. These findings show that histogram relation learning enhances the diversity of the latent space. We then incorporate the residual learning module, which improves the rFID from 0.97 to 0.92 and improves the PSNR from 22.30 to 22.44. These results demonstrate the effectiveness of the residual learning module in optimizing the reconstruction performance of the tokenizer.

Histogram Relation	Residual	rFID ↓	PSNR ↑
×	×	1.43	22.06
✓	×	0.97	22.30
✓	✓	0.92	22.44

Table 5: Ablation study on different components for the improvement of GloTok.

5 Conclusion

In this paper, we introduce GloTok, an image tokenizer that leverages global relational information to achieve a more uniform semantic latent distribution. Employing a codebook-wise histogram relation learning method, GloTok effectively regulates the global distribution of tokens across the entire dataset. GloTok does not have to directly access the pre-trained model during training, effectively reducing both training time and GPU memory usage. In addition, a residual learning module is proposed to mitigate the fine-grained detail lost during discretization. Experiments on the ImageNet dataset show that GloTok achieves state-of-the-art reconstruction performance and generative performance.

Acknowledgments

The computations in this research involved with Yuxi Mi and Shuigeng Zhou were performed using the CFFF platform of Fudan University.

References

- Arora, K.; Asri, L. E.; Bahuleyan, H.; and Cheung, J. C. K. 2023. Why Exposure Bias Matters: An Imitation Learning Perspective of Error Accumulation in Language Generation. *arXiv:2204.01171*.
- Bai, Z.; Gao, J.; Gao, Z.; Wang, P.; Zhang, Z.; He, T.; and Shou, M. Z. 2024. Factorized Visual Tokenization and Generation. *arXiv preprint arXiv:2411.16681*.
- Cao, S.; Yin, Y.; Huang, L.; Liu, Y.; Zhao, X.; Zhao, D.; and Huang, K. 2023. Efficient-vqgan: Towards high-resolution image generation with efficient vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7368–7377.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging Properties in Self-Supervised Vision Transformers. *CoRR*, abs/2104.14294.
- Darcet, T.; Oquab, M.; Mairal, J.; and Bojanowski, P. 2023. Vision Transformers Need Registers.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12873–12883.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; Klambauer, G.; and Hochreiter, S. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Nash Equilibrium. *CoRR*, abs/1706.08500.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1125–1134.
- Kingma, D. P.; and Welling, M. 2022. Auto-Encoding Variational Bayes. *arXiv:1312.6114*.
- Lee, D.; Kim, C.; Kim, S.; Cho, M.; and Han, W.-S. 2022. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11523–11532.
- Li, T.; Tian, Y.; Li, H.; Deng, M.; and He, K. 2024a. Autoregressive Image Generation without Vector Quantization. *arXiv:2406.11838*.
- Li, X.; Qiu, K.; Chen, H.; Kuen, J.; Gu, J.; Raj, B.; and Lin, Z. 2024b. Imagefolder: Autoregressive image generation with folded tokens. *arXiv preprint arXiv:2410.01756*.
- Luo, Z.; Shi, F.; Ge, Y.; Yang, Y.; Wang, L.; and Shan, Y. 2024. Open-magvit2: An open-source project toward democratizing auto-regressive visual generation. *arXiv preprint arXiv:2409.04410*.
- Ma, C.; Jiang, Y.; Wu, J.; Yang, J.; Yu, X.; Yuan, Z.; Peng, B.; and Qi, X. 2025a. UniTok: A Unified Tokenizer for Visual Generation and Understanding. *arXiv preprint arXiv:2502.20321*.
- Ma, C.; Jiang, Y.; Wu, J.; Yang, J.; Yu, X.; Yuan, Z.; Peng, B.; and Qi, X. 2025b. UniTok: A Unified Tokenizer for Visual Generation and Understanding. *arXiv preprint arXiv:2502.20321*.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H. V.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; Howes, R.; Huang, P.-Y.; Xu, H.; Sharma, V.; Li, S.-W.; Galuba, W.; Rabbat, M.; Assran, M.; Ballas, N.; Synnaeve, G.; Misra, I.; Jégou, H.; Mairal, J.; Labatut, P.; Joulin, A.; and Bojanowski, P. 2023. DINOv2: Learning Robust Visual Features without Supervision.
- Qu, L.; Zhang, H.; Liu, Y.; Wang, X.; Jiang, Y.; Gao, Y.; Ye, H.; Du, D. K.; Yuan, Z.; and Wu, X. 2024. Tokenflow: Unified image tokenizer for multimodal understanding and generation. *arXiv preprint arXiv:2412.03069*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training.
- Ren, S.; Yu, Q.; He, J.; Shen, X.; Yuille, A.; and Chen, L.-C. 2024. FlowAR: Scale-wise Autoregressive Image Generation Meets Flow Matching. *arXiv preprint arXiv:2412.15205*.
- Ren, S.; Yu, Q.; He, J.; Shen, X.; Yuille, A.; and Chen, L.-C. 2025. Beyond Next-Token: Next-X Prediction for Autoregressive Visual Generation. *arXiv:2502.20388*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved techniques for training gans. *Advances in neural information processing systems*, 29.
- Shi, F.; Luo, Z.; Ge, Y.; Yang, Y.; Shan, Y.; and Wang, L. 2024. Taming scalable visual tokenizer for autoregressive image generation. *arXiv preprint arXiv:2412.02692*.

Sun, P.; Jiang, Y.; Chen, S.; Zhang, S.; Peng, B.; Luo, P.; and Yuan, Z. 2024. Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint arXiv:2406.06525*.

Tian, K.; Jiang, Y.; Yuan, Z.; Peng, B.; and Wang, L. 2024. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *Advances in neural information processing systems*, 37: 84839–84865.

Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention Is All You Need. *CoRR*, abs/1706.03762.

Weber, M.; Yu, L.; Yu, Q.; Deng, X.; Shen, X.; Cremers, D.; and Chen, L.-C. 2024. MaskBit: Embedding-free Image Generation via Bit Tokens. *arXiv:2409.16211*.

Yao, J.; Yang, B.; and Wang, X. 2025. Reconstruction vs. generation: Taming optimization dilemma in latent diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 15703–15712.

Yu, J.; Li, X.; Koh, J. Y.; Zhang, H.; Pang, R.; Qin, J.; Ku, A.; Xu, Y.; Baldrige, J.; and Wu, Y. 2021. Vector-quantized image modeling with improved vqgan. *arXiv preprint arXiv:2110.04627*.

Yu, L.; Lezama, J.; Gundavarapu, N. B.; Versari, L.; Sohn, K.; Minnen, D.; Cheng, Y.; Birodkar, V.; Gupta, A.; Gu, X.; et al. 2023. Language Model Beats Diffusion–Tokenizer is Key to Visual Generation. *arXiv preprint arXiv:2310.05737*.

Yu, Q.; He, J.; Deng, X.; Shen, X.; and Chen, L.-C. 2024a. Randomized autoregressive visual generation. *arXiv preprint arXiv:2411.00776*.

Yu, Q.; Weber, M.; Deng, X.; Shen, X.; Cremers, D.; and Chen, L.-C. 2024b. An image is worth 32 tokens for reconstruction and generation. *Advances in Neural Information Processing Systems*, 37: 128940–128966.

Zhu, L.; Wei, F.; Lu, Y.; and Chen, D. 2024a. Scaling the codebook size of vqgan to 100,000 with a utilization rate of 99%. *arXiv preprint arXiv:2406.11837*.

Zhu, Y.; Li, B.; Xin, Y.; and Xu, L. 2024b. Addressing representation collapse in vector quantized models with one linear layer. *arXiv preprint arXiv:2411.02038*.

Zhu, Y.; Li, B.; Zhang, H.; Li, X.; Xu, L.; and Bing, L. 2024c. Stabilize the latent space for image autoregressive modeling: A unified perspective. *arXiv preprint arXiv:2410.12490*.